arXiv:1610.02967v1 [math.OC] 7 Oct 2016

Distributed Convex Optimization with Many Convex Constraints

Joachim Giesen and Sören Laue, Friedrich-Schiller-University Jena, Germany
{joachim.giesen,soeren.laue}@uni-jena.de

Abstract

We address the problem of solving convex optimization problems with many convex constraints in a distributed setting. Our approach is based on an extension of the alternating direction method of multipliers (ADMM) that recently gained a lot of attention in the Big Data context. Although it has been invented decades ago, ADMM so far can be applied only to unconstrained problems and problems with linear equality or inequality constraints. Our extension can handle arbitrary inequality constraints directly. It combines the ability of ADMM to solve convex optimization problems in a distributed setting with the ability of the Augmented Lagrangian method to solve constrained optimization problems, and as we show, it inherits the convergence guarantees of ADMM and the Augmented Lagrangian method.

1 Introduction

An optimization problem can be called large scale if it involves a large number of optimization variables, and/or a large number of input parameters, and/or a large number of constraints. The increasing availability of distributed hardware suggests to address large scale optimization problems by distributed algorithms. Here we design and analyze a distributed algorithm for general convex optimization problems that involve a large number of constraints. The algorithm is a variant of the alternating direction method of multipliers (ADMM) that was proposed by Glowinski and Marroco [7] and by Gabay and Mercier [4].

Recently, ADMM regained significant attention, especially in the machine learning community, because it allows to solve convex optimization problems that involve a large number of parameters in a distributed setting [2]. In a typical machine learning method, like the least squares method for regression problems, the parameters are just the data points. Hence, ADMM is often the method of choice for Big Data problems, where the data does not fit into the memory of a single compute node. The optimization problem behind a typical Big Data machine learning method usually aims at minimizing a so called loss-function that is the sum of the losses for each data point. Hence, the objective function f of such problems is separable, i.e., it holds that $f(x) = \sum_i f_i(x_i)$, where f_i is determined by the i-th data point. In this case ADMM lends itself to a distributed implementation where the data points are distributed.

ADMM works for unconstrained optimization problems and for optimization problems with linear equality and/or inequality constraints. Surprisingly, so far no general convex inequality constraints have been considered directly in the context of ADMM. Here we extend ADMM such that it can deal with arbitrary convex inequality constraints. The extension can handle a large number of constraints by distributing them.

Optimization problems with a large number of constraints typically also arise as Big Data problems, but instead of contributing a term to the objective function each data point now contributes a constraint to the problem. An illustrative example from computational geometry is the smallest enclosing ball problem [3, 6, 14], where we are given data points in Euclidean space and have to compute the smallest ball that contains all the points, i.e., its radius and center. The objective function here is the radius of the ball that needs to be minimized, and every data point contributes a constraint, namely the distance of the point from the center must be at most the radius. The smallest enclosing ball problem can be solved by our extension of ADMM in a distributed setting.

Related work. To the best of our knowledge our extension of ADMM is the first distributed algorithm to solve general convex optimization problems with no restrictions on the type of constraints or assumptions on the structure of the problem.

Mosk-Aoyama et al. [12] have designed and analyzed a distributed algorithm for solving convex optimization problems with separable objective function and linear equality constraints. Their algorithm blends a gossip-based information spreading, iterative gradient ascent method with the barrier method from interior-point algorithms. It is similar to ADMM and can also handle only linear constraints.

Zhu and Martínez [16] have introduced a distributed multi-agent algorithm for minimizing a convex function that is the sum of local functions subject to a global equality or inequality constraint. Their algorithm involves a projection onto local constraint sets which are usually as hard to compute as solving the original problem with general constraints. For instance, it is well known via standard duality theory that the feasibility problem for linear programs is as hard as solving linear programs. This holds true for general convex optimization problems with vanishing duality gap.

In principle, the standard ADMM can also handle convex constraints by transforming them into indicator functions that are added to the objective function. However, this leads to subproblems that ADMM needs to solve in each iteration that entail computing a projection onto the feasible region. This leads to the same issues as with the method by Zhu and Martínes [16] since computing these projections can be as hard as solving the original problem. We will elaborate on this in more detail in Section 2.

There are also quite a few parallel or distributed algorithms for specific optimization problems like packing and covering linear programs. We discuss some of them in Section 6.

Paper outline. In Section 2 we review the alternating direction method of multipliers (ADMM) and briefly discuss how it can be adapted to deal with Big Data problems in a distributed setting. Our goal is to extend ADMM such that it can also handle (many) convex inequality constraints and that these constraints can be distributed similarly as the data is distributed in the standard ADMM. In Section 3 we state the class of convex optimization problems that we want to solve with our extension of ADMM and describe the extension itself. Section 4 is devoted to a convergence analysis of our extension of ADMM. In Section 5 we use our extension of ADMM for solving convex optimization problems with a large number of convex equality and inequality constraints in a distributed setting. The paper is then concluded in Section 6 with a brief discussion of potential applications.

2 Alternating direction method of multipliers (ADMM)

ADMM is an iterative algorithm that in its most general form can solve convex optimization problems of the form

$$\begin{array}{ll} \min_{x,z} & f_1(x) + f_2(z) \\ \text{s.t.} & Ax + Bz - c = 0, \end{array} \tag{1}$$

where $f_1 : \mathbb{R}^{n_1} \to \mathbb{R} \cup \{\infty\}$ and $f_2 : \mathbb{R}^{n_2} \to \mathbb{R} \cup \{\infty\}$ are convex functions, $A \in \mathbb{R}^{m \times n_1}$ and $B \in \mathbb{R}^{m \times n_2}$ are matrices, and $c \in \mathbb{R}^m$.

ADMM can obviously deal with linear equality constraints, but it can also handle linear inequality constraints. The latter are reduced to linear equality constraints by replacing constraints of the form $Ax \leq b$ by $Ax + s = b$, adding the slack variable s to the set of optimization variables, and setting $f_2(s) = \mathbb{1}_{\mathbb{R}^n_+}(s)$, where

$$\mathbb{1}_{\mathbb{R}^n_+}(s) = \begin{cases} 0, & \text{if } s \geq 0 \\ \infty, & \text{otherwise,} \end{cases}$$

is the indicator function of the set $\mathbb{R}^n_+ = \{x \in \mathbb{R}^n | x \geq 0\}$. Note that f_1 and f_2 are allowed to take the value ∞.

Recently, ADMM regained a lot of attention, because it allows to solve Big Data problems with separable objective function in a distributed setting. Such problems are typically given as

$$\min_x \quad \sum_i f_i(x),$$

where f_i corresponds to the i-th data point (or more generally i-th data block) and x is a weight vector that describes the data model. This problem can be transformed into the equivalent optimization problem, with individual weight vectors x_i for each data point (data block) that are coupled through an equality constraint,

$$\begin{array}{ll} \min_{x_i,z} & \sum_i f_i(x_i) \\ \text{s.t.} & x_i - z = 0 \quad \forall\, i, \end{array}$$

which is a special case of Problem 1 that can be solved by ADMM in a distributed setting by distributing the data.

Adding convex inequality constraints to Problem 1 does not destroy convexity of the problem, but so far ADMM cannot deal with such constraints. Note that the problem only remains convex, if all equality constraints are induced by affine functions. That is, we cannot add convex equality constraints in general without destroying convexity.

Our goal for the following sections is to extend ADMM such that it can also deal with nonlinear, convex inequality constraints. For problems with many constraints we will show that these constraints can be distributed similarly as the data points in Big Data problems with separable objective function are distributed for the standard ADMM.

In the introduction we have already mentioned an immediate idea for reaching this goal, namely the constraints can be transformed into indicator functions that are then added to the objective function. Given a set of constraints $g(x) \leq 0$, adding an indicator function $\mathbb{1}_{\{g(x) \leq 0\}}$ to the objective function allows to remove the constraints. However, ADMM needs to compute a projection onto the feasible set $\{x | g(x) \leq 0\}$ in every iteration. Since computing these projections can be as hard as solving the original problem we have just deferred the

difficulties to the subproblems. Only in special cases, for instance when the constraint set is the positive orthant, the projections can be carried out more efficiently. However, in general this is not the case. Here, we will devise a method for dealing with arbitrary constraints directly without any hard-to-compute projections.

3 ADMM for problems with convex constraints

Here we introduce our extension of ADMM that can also solve convex optimization problems with convex inequality constraints. That is, we consider convex optimization problems of the form

$$\begin{array}{ll} \min_{x,z} & f_1(x) + f_2(z) \\ \text{s.t.} & g_0(x) \leq 0 \\ & h_1(x) + h_2(z) = 0, \end{array} \tag{2}$$

where f_1 and f_2 are as in Problem 1, $g_0 : \mathbb{R}^{n_1} \to \mathbb{R}^p$ is convex in every component, and $h_1 : \mathbb{R}^{n_1} \to \mathbb{R}^m$ and $h_2 : \mathbb{R}^{n_2} \to \mathbb{R}^m$ are affine functions. In the following we assume that the problem is feasible, i.e., that a feasible solution exists, and that strong duality holds. A sufficient condition for strong duality is that the interior of the feasible region is non-empty. This condition is known as Slater's condition for convex optimization problems.

For our extension of ADMM and its convergence analysis we need to work with an equivalent reformulation of Problem 2, where we replace $g_0(x)$ by $g(x) = \max\{0, g_0(x)\}$, with componentwise maximum, and turn the convex inequality constraints into convex equality constraints. Thus, in the following we are considering optimization problems of the form

$$\begin{array}{ll} \min_{x,z} & f_1(x) + f_2(z) \\ \text{s.t.} & g(x) = 0 \\ & h_1(x) + h_2(z) = 0, \end{array} \tag{3}$$

where $g(x) = \max\{0, g_0(x)\}$, which by construction is again convex in every component. Note though that Problem 3 is no longer convex, because $g(x) = 0$ is not an affine constraint. However, we show in the following that it can be solved efficiently.

Analogously to ADMM our extension builds on the the Augmented Lagrangian for Problem 3 which is the following function

$$L_\rho(x, z, \mu, \lambda) = f_1(x) + f_2(z) + \frac{\rho}{2} \|g(x)\|^2 + \mu^\top g(x) + \frac{\rho}{2} \|h_1(x) + h_2(z)\|^2 + \lambda^\top (h_1(x) + h_2(z)),$$

where $\mu \in \mathbb{R}^p$ and $\lambda \in \mathbb{R}^m$ are Lagrange multipliers, $\rho > 0$ is some constant, and $\|\cdot\|$ denotes the Euclidean norm. The Lagrange multipliers are also referred to as dual variables.

Algorithm 1 is our extension of ADMM for solving instances of Problem 3. It runs in iterations. In the $(k+1)$-th iteration the primal variables x^k and z^k as well as the dual variables μ^k and λ^k are updated.

4 Convergence analysis

From duality theory we know that for all $x \in \mathbb{R}^{n_1}$ and $z \in \mathbb{R}^{n_2}$

$$L_0(x^*, z^*, \mu^*, \lambda^*) \leq L_0(x, z, \mu^*, \lambda^*), \tag{4}$$

Algorithm 1 ADMM for problems with convex inequality constraints

1: **input:** instance of Problem 3
2: **output:** approximate solution $x \in \mathbb{R}^{n_1}, z \in \mathbb{R}^{n_2}, \mu \in \mathbb{R}^p, \lambda \in \mathbb{R}^m$
3: initialize $x^0 = 0$, $z^0 = 0$, $\mu^0 = 0$, $\lambda^0 = 0$, and ρ to some constant > 0
4: **repeat**
5: $\quad x^{k+1} := \quad \operatorname{argmin}_x \; L_\rho(x, z^k, \mu^k, \lambda^k)$
6: $\quad z^{k+1} := \quad \operatorname{argmin}_z \; L_\rho(x^{k+1}, z, \mu^k, \lambda^k)$
7: $\quad \mu^{k+1} := \quad \mu^k + \rho g(x^{k+1})$
8: $\quad \lambda^{k+1} := \quad \lambda^k + \rho \left(h_1(x^{k+1}) + h_2(z^{k+1})\right)$
9: **until** convergence
10: **return** $x^k, z^k, \mu^k, \lambda^k$

where L_0 is the Lagrangian of Problem 3 and x^*, z^*, μ^*, and λ^* are optimal primal and dual variables. Note, that x^*, z^*, μ^*, and λ^* are not necessarily unique. Here, they refer to just one optimal solution. Also note that the Lagrangian is identical to the Augmented Lagrangian with $\rho = 0$. Given that strong duality holds the optimal solution to the original Problem 3 is identical to the optimal solution of the Lagrangian, i.e., the dual.

We need a few more definitions. Let $f^k = f_1(x^k) + f_2(z^k)$ be the objective function value at the k-th iterate (x^k, z^k) and let f^* be the optimal function value. Let $r_g^k = g(x^k)$ be the residual of the nonlinear equality constraints, i.e., the constraints originating from the convex inequality constraints, and let $r_h^k = h_1(x^k) + h_2(z^k)$ be the residual of the linear equality constraints in iteration k.

Our goal in this section is to prove the following theorem.

Theorem 1. *When Algorithm 1 is applied to an instance of Problem 3, then*

$$\lim_{k\to\infty} r_g^k = 0, \quad \lim_{k\to\infty} r_h^k = 0, \quad \textit{and} \quad \lim_{k\to\infty} f^k = f^*.$$

The theorem states primal feasibility and convergence of the primal objective function value. Note, however, that convergence to primal optimal points x^* and z^* cannot be guaranteed. This is the case for the original ADMM as well. Additional assumptions on the functions f_1, f_2, and g, like, for instance, strong convexity, are necessary to guarantee convergence to the primal optimal points. However, the points x^k, z^k will be primal optimal and feasible up to an arbitrarily small error for sufficiently large k.

The proof of Theorem 1 follows along the lines of the convergence proof for the original ADMM in [2] and is subdivided into four lemmas.

Lemma 1. *The dual variables μ^k are non-negative for all iterations, i.e., it holds that $\mu^k \geq 0$ for all $k \in \mathbb{N}$.*

Proof. The proof is by induction. In Line 3 of Algorithm 1 the dual variable is initialized as $\mu^0 = 0$. If $\mu^k \geq 0$, then it follows from the update rule in Line 7 of Algorithm 1 that $\mu^{k+1} = \mu^k + \rho g(x) \geq 0$, since $g(x) = \max\{0, g_0(x)\} \geq 0$ and by assumption also $\rho > 0$. □

Lemma 2. *The difference between the optimal objective function value and its value at the $(k+1)$-th iterate can be bounded as*

$$f^* - f^{k+1} \leq (\mu^*)^\top r_g^{k+1} + (\lambda^*)^\top r_h^{k+1}$$

Proof. It follows from the definitions, vanishing constraints in an optimum, and Inequality 4 that

$$\begin{aligned} f^* &= f_1(x^*) + f_2(z^*) \\ &= L_0(x^*, z^*, \mu^*, \lambda^*) \\ &\leq L_0(x^{k+1}, z^{k+1}, \mu^*, \lambda^*) \\ &= f_1(x^{k+1}) + f_2(z^{k+1}) + (\mu^*)^\top r_g^{k+1} + (\lambda^*)^\top r_h^{k+1} \\ &= f^{k+1} + (\mu^*)^\top r_g^{k+1} + (\lambda^*)^\top r_h^{k+1}. \end{aligned}$$

□

Lemma 3. *The difference between the value of the objective function at the $(k+1)$-th iterate and its optimal value can be bounded as follows*

$$f^{k+1} - f^* \leq -(\mu^{k+1})^\top r_g^{k+1} - (\lambda^{k+1})^\top r_h^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(-r_h^{k+1} + h_2(z^{k+1}) - h_2(z^*)\right).$$

Proof. From Line 5 of Algorithm 1 we know that x^{k+1} minimizes function $L_\rho(x, z^k, \mu^k, \lambda^k)$ with respect to x. Hence, we know that 0 must be contained in the subdifferential of $L_\rho(x, z^k, \mu^k, \lambda^k)$ with respect to x at x^{k+1}, i.e.,

$$\begin{aligned} 0 \in\ & \partial f_1(x^{k+1}) + \rho \cdot \partial g(x^{k+1}) \cdot g(x^{k+1}) + \partial g(x^{k+1}) \cdot \mu^k \\ & + \rho \cdot \partial h_1(x^{k+1}) \cdot \left(h_1(x^{k+1}) + h_2(z^k)\right) + \partial h_1(x^{k+1}) \cdot \lambda^k, \end{aligned}$$

where $\partial f_1(x^{k+1}) \in \mathbb{R}^{n_1}$ is the subdifferential of f_1 at x^{k+1}, $\partial g(x^{k+1}) \in \mathbb{R}^{n_1 \times p}$ is the subdifferential of g at x^{k+1}, and $\partial h_1(x^{k+1}) \in \mathbb{R}^{n_1 \times m}$ is the subdifferential of h_1 at x^{k+1}.

The update rule for the dual variables μ in Line 7 of Algorithm 1 gives

$$\mu^k = \mu^{k+1} - \rho g(x^{k+1})$$

and similarly, the update rule for the dual variables λ in Line 8 gives

$$\lambda^k = \lambda^{k+1} - \rho \left(h_1(x^{k+1}) + h_2(z^{k+1})\right).$$

Plugging these update rules into the subdifferential optimality condition from above gives

$$\begin{aligned} 0 \in\ & \partial f_1(x^{k+1}) + \rho \cdot \partial g(x^{k+1}) \cdot g(x^{k+1}) + \partial g(x^{k+1}) \cdot \left(\mu^{k+1} - \rho g(x^{k+1})\right) \\ & + \rho \cdot \partial h_1(x^{k+1}) \cdot \left(h_1(x^{k+1}) + h_2(z^k)\right) + \partial h_1(x^{k+1}) \cdot \left(\lambda^{k+1} - \rho h_1(x^{k+1}) - \rho h_2(z^{k+1})\right), \end{aligned}$$

and thus

$$0 \in \partial f_1(x^{k+1}) + \partial g(x^{k+1}) \cdot \mu^{k+1} + \partial h_1(x^{k+1}) \cdot \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right).$$

If 0 is contained in the subdifferential of a convex function at point x, then x is a minimizer of this function. That is, x^{k+1} minimizes the convex function

$$x \mapsto f_1(x) + (g(x))^\top \mu^{k+1} + (h_1(x))^\top \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right). \tag{5}$$

This function is convex, because f_1 and g are convex functions, h_1 is an affine function, and any non-negative combination of convex functions is again a convex function. Note that we have $\mu^{k+1} \geq 0$ by Lemma 1.

Similarly, Line 6 of Algorithm 1 implies that 0 is contained in the subdifferential of $L_\rho(x^{k+1}, z, \mu^k, \lambda^k)$ with respect to z at z^{k+1}, i.e.,

$$0 \in \partial f_2(z^{k+1}) + \rho \cdot \partial h_2(z^{k+1}) \cdot \left(h_1(x^{k+1}) + h_2(z^{k+1})\right) + \partial h_2(z^{k+1}) \cdot \lambda^k.$$

Again, substituting $\lambda^k = \lambda^{k+1} - \rho \left(h_1(x^{k+1}) + h_2(z^{k+1})\right)$ we get

$$0 \in \partial f_2(z^{k+1}) + \partial h_2(z^{k+1}) \cdot \lambda^{k+1}.$$

Hence, z^{k+1} minimizes the convex function

$$z \mapsto f_2(z) + (\lambda^{k+1})^\top h_2(z). \tag{6}$$

The function is convex since f_2 is convex and h_2 is affine.

Since x^{k+1} is a minimizer of Function 5 we have

$$\begin{aligned} &f_1(x^{k+1}) + (g(x^{k+1}))^\top \mu^{k+1} + (h_1(x^{k+1}))^\top \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right) \\ &\leq\ f_1(x^*) + (g(x^*))^\top \mu^{k+1} + (h_1(x^*))^\top \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right). \end{aligned}$$

Analogously, since z^{k+1} minimizes Function 6 we have

$$f_2(z^{k+1}) + (\lambda^{k+1})^\top h_2(z^{k+1}) \leq f_2(z^*) + (\lambda^{k+1})^\top h_2(z^*).$$

Finally, from summing up both inequalities and rearranging we get

$$\begin{aligned} f^{k+1} - f^* \ &= \ f_1(x^{k+1}) + f_2(z^{k+1}) - f_1(x^*) - f_2(z^*) \\ &\leq (g(x^*))^\top \mu^{k+1} + (h_1(x^*))^\top \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right) - (g(x^{k+1}))^\top \mu^{k+1} \\ &\quad - (h_1(x^{k+1}))^\top \left(\lambda^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)\right) + (\lambda^{k+1})^\top h_2(z^*) - (\lambda^{k+1})^\top h_2(z^{k+1}) \\ &= -(g(x^{k+1}))^\top \mu^{k+1} + (\lambda^{k+1})^\top \left(h_1(x^*) + h_2(z^*)\right) - (\lambda^{k+1})^\top \left(h_1(x^{k+1}) + h_2(z^{k+1})\right) \\ &\quad - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(h_1(x^*) - h_1(x^{k+1})\right) \\ &= -(\mu^{k+1})^\top r_g^{k+1} - (\lambda^{k+1})^\top r_h^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(h_1(x^*) - h_1(x^{k+1})\right) \\ &= -(\mu^{k+1})^\top r_g^{k+1} - (\lambda^{k+1})^\top r_h^{k+1} \\ &\quad - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(h_1(x^*) + h_2(z^*) - h_2(z^*) - h_1(x^{k+1}) - h_2(z^{k+1}) + h_2(z^{k+1})\right) \\ &= -(\mu^{k+1})^\top r_g^{k+1} - (\lambda^{k+1})^\top r_h^{k+1} - \rho \left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(-h_2(z^*) - r_h^{k+1} + h_2(z^{k+1})\right), \end{aligned}$$

where we have used that $g(x^*) = 0$, $h_1(x^*) + h_2(z^*) = 0$, $r_g^{k+1} = g(x^{k+1})$ is the residual of the convex constraints, and $r_h^{k+1} = h_1(x^{k+1}) + h_2(z^{k+1})$ is the residual of the affine constraints in iteration $k+1$. □

To continue, we need one more definition.

Definition 1. *Let* $V^k = \frac{1}{\rho}\|\mu^k - \mu^*\|^2 + \frac{1}{\rho}\|\lambda^k - \lambda^*\|^2 + \rho\|h_2(z^k) - h_2(z^*)\|^2$.

For this newly defined quantity we show in the following lemma that it is non-increasing over the iterations. This property will be crucial in the proof of Theorem 1.

Lemma 4. *For every iteration* $k \in \mathbb{N}$ *it holds that*

$$V^k - V^{k+1} \geq \rho\|r_g^{k+1}\|^2 + \rho\|r_h^{k+1}\|^2 + \rho\|h_2(z^{k+1}) - h_2(z^k)\|^2.$$

Proof of Lemma 4. Summing up the inequality in Lemma 2 and the inequality in Lemma 3 gives

$$\begin{aligned} 0 \leq\ & (\mu^*)^\top r_g^{k+1} - (\mu^{k+1})^\top r_g^{k+1} + (\lambda^*)^\top r_h^{k+1} - (\lambda^{k+1})^\top r_h^{k+1} \\ & - \rho\left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(-r_h^{k+1} + h_2(z^{k+1}) - h_2(z^*)\right), \end{aligned}$$

or equivalently, by rearranging and multiplying by 2,

$$\begin{aligned} 0 \geq\ & 2(\mu^{k+1} - \mu^*)^\top r_g^{k+1} + 2(\lambda^{k+1} - \lambda^*)^\top r_h^{k+1} \\ & + 2\rho\left(h_2(z^{k+1}) - h_2(z^k)\right)^\top \left(-r_h^{k+1} + h_2(z^{k+1}) - h_2(z^*)\right). \end{aligned} \tag{7}$$

Next we are rewriting the three terms in this inequality individually.

Using the update rule $\mu^{k+1} = \mu^k + \rho g(x^{k+1}) = \mu^k + \rho r_g^{k+1}$ for the Lagrange multipliers μ in Line 7 of Algorithm 1 several times we can rewrite the first term as follows

$$\begin{aligned} & 2(\mu^{k+1} - \mu^*)^\top r_g^{k+1} \\ &= 2(\mu^k + \rho r_g^{k+1} - \mu^*)^\top r_g^{k+1} \\ &= 2(\mu^k - \mu^*)^\top r_g^{k+1} + \rho\|r_g^{k+1}\|^2 + \rho\|r_g^{k+1}\|^2 \\ &= 2\frac{(\mu^k - \mu^*)^\top(\mu^{k+1} - \mu^k)}{\rho} + \frac{\|\mu^{k+1} - \mu^k\|^2}{\rho} + \rho\|r_g^{k+1}\|^2 \\ &= 2\frac{(\mu^k - \mu^*)^\top\left(\mu^{k+1} - \mu^* - (\mu^k - \mu^*)\right)}{\rho} \\ &\quad + \frac{\|\mu^{k+1} - \mu^* - (\mu^k - \mu^*)\|^2}{\rho} + \rho\|r_g^{k+1}\|^2 \\ &= \frac{2(\mu^k - \mu^*)^\top\left(\mu^{k+1} - \mu^*\right) - 2\|\mu^k - \mu^*\|^2}{\rho} \\ &\quad + \frac{\|\mu^{k+1} - \mu^*\|^2 + \|\mu^k - \mu^*\|^2 - 2(\mu^{k+1} - \mu^*)^\top(\mu^k - \mu^*)}{\rho} + \rho\|r_g^{k+1}\|^2 \\ &= \frac{\|\mu^{k+1} - \mu^*\|^2 - \|\mu^k - \mu^*\|^2}{\rho} + \rho\|r_g^{k+1}\|^2. \end{aligned}$$

The analogous argument holds for the second term, when using the update rule $\lambda^{k+1} = \lambda^k + \rho r_h^{k+1}$ in Line 8 of Algorithm 1, i.e., we have

$$2(\lambda^{k+1} - \lambda^*)^\top r_h^{k+1} = \frac{\|\lambda^{k+1} - \lambda^*\|^2 - \|\lambda^k - \lambda^*\|^2}{\rho} + \rho\|r_h^{k+1}\|^2.$$

Adding $\rho \left\| r_h^{k+1} \right\|^2$ to the third term of Inequality 7 gives

$$
\begin{aligned}
&\rho \left\| r_h^{k+1} \right\|^2 + 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top \left(-r_h^{k+1} + h_2(z^{k+1}) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} \|^2 - 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top r_h^{k+1} + 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top \left(h_2(z^{k+1}) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} \|^2 - 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top r_h^{k+1} \\
&\quad + 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top \left(\left(h_2(z^{k+1}) - h_2(z^k) \right) + \left(h_2(z^k) - h_2(z^*) \right) \right) \\
&= \rho \| r_h^{k+1} \|^2 - 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top r_h^{k+1} \\
&\quad + 2\rho \| h_2(z^{k+1}) - h_2(z^k) \|^2 + 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} - \left(h_2(z^{k+1}) - h_2(z^k) \right) \|^2 + \rho \| h_2(z^{k+1}) - h_2(z^k) \|^2 \\
&\quad + 2\rho \left(h_2(z^{k+1}) - h_2(z^k) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} - \left(h_2(z^{k+1}) - h_2(z^k) \right) \|^2 + \rho \| \left(h_2(z^{k+1}) - h_2(z^*) \right) - \left(h_2(z^k) - h_2(z^*) \right) \|^2 \\
&\quad + 2\rho \left(\left(h_2(z^{k+1}) - h_2(z^*) \right) - \left(h_2(z^k) - h_2(z^*) \right) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} - \left(h_2(z^{k+1}) - h_2(z^k) \right) \|^2 + \rho \| h_2(z^{k+1}) - h_2(z^*) \|^2 + \rho \| h_2(z^k) - h_2(z^*) \|^2 \\
&\quad - 2\rho \left(h_2(z^{k+1}) - h_2(z^*) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) \\
&\quad + 2\rho \left(h_2(z^{k+1}) - h_2(z^*) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) - 2\rho \left(h_2(z^k) - h_2(z^*) \right)^\top \left(h_2(z^k) - h_2(z^*) \right) \\
&= \rho \| r_h^{k+1} - \left(h_2(z^{k+1}) - h_2(z^k) \right) \|^2 + \rho \| h_2(z^{k+1}) - h_2(z^*) \|^2 - \rho \| h_2(z^k) - h_2(z^*) \|^2.
\end{aligned}
$$

Hence, Inequality 7 is equivalent to

$$
\begin{aligned}
0 \geq\ & \frac{1}{\rho} \left(\| (\mu^{k+1} - \mu^* \|^2 - \| \mu^k - \mu^* \|^2 \right) + \rho \| r_g^{k+1} \|^2 \\
& + \frac{1}{\rho} \left(\| (\lambda^{k+1} - \lambda^* \|^2 - \| \lambda^k - \lambda^* \|^2 \right) \\
& + \rho \| r_h^{k+1} - \left(h_2(z^{k+1}) - h_2(z^k) \right) \|^2 + \rho \| h_2(z^{k+1}) - h_2(z^*) \|^2 - \rho \| h_2(z^k) - h_2(z^*) \|^2.
\end{aligned}
$$

By rearranging the terms in this inequality and using the following term expansion

$$
\| r_h^{k+1} - (h_2(z^{k+1}) - h_2(z^k)) \|^2 = \| r_h^{k+1} \|^2 + \| h_2(z^{k+1}) - h_2(z^k) \|^2 - 2(r_h^{k+1})^\top \left(h_2(z^{k+1}) - h_2(z^k) \right),
$$

we get

$$\begin{aligned}
&\rho\|r_g^{k+1}\|^2 + \rho\|r_h^{k+1}\|^2 + \rho\|h_2(z^{k+1}) - h_2(z^k)\|^2 - 2\rho(r_h^{k+1})^\top \Big(h_2(z^{k+1}) - h_2(z^k)\Big) \\
&\quad = \rho\|r_g^{k+1}\|^2 + \rho\|r_h^{k+1} - \Big(h_2(z^{k+1}) - h_2(z^k)\Big)\|^2 \\
&\quad \leq \frac{1}{\rho}\|\mu^k - \mu^*\|^2 + \frac{1}{\rho}\|\lambda^k - \lambda^*\|^2 + \rho\|h_2(z^k) - h_2(z^*)\|^2 \\
&\qquad - \left(\frac{1}{\rho}\|\mu^{k+1} - \mu^*\|^2 + \frac{1}{\rho}\|\lambda^{k+1} - \lambda^*\|^2 + \rho\|h_2(z^{k+1}) - h_2(z^*)\|^2\right) \\
&\quad = V^k - V^{k+1},
\end{aligned}$$

where we have used Definition 1 of V^k in the last equality. Hence, to finish the proof of Lemma 4 it only remains to show that

$$2\rho(r_h^{k+1})^\top \Big(h_2(z^{k+1}) - h_2(z^k)\Big) \leq 0.$$

From the proof of Lemma 3 we know that z^{k+1} minimizes the function $f_2(z) + (\lambda^{k+1})^\top h_2(z)$ and similarly that z^k minimizes the function $f_2(z) + (\lambda^k)^\top h_2(z)$. Hence, we have the following two inequalities

$$f_2(z^{k+1}) + (\lambda^{k+1})^\top h_2(z^{k+1}) \leq f_2(z^k) + (\lambda^{k+1})^\top h_2(z^k)$$

and

$$f_2(z^k) + (\lambda^k)^\top h_2(z^k) \leq f_2(z^{k+1}) + (\lambda^k)^\top h_2(z^{k+1}).$$

Summing up these two inequalities yields

$$(\lambda^{k+1})^\top h_2(z^{k+1}) + (\lambda^k)^\top h_2(z^k) \leq (\lambda^{k+1})^\top h_2(z^k) + (\lambda^k)^\top h_2(z^{k+1}),$$

or equivalently

$$\begin{aligned}
0 \;&\geq\; (\lambda^{k+1})^\top \Big(h_2(z^{k+1}) - h_2(z^k)\Big) + (\lambda^k)^\top \Big(h_2(z^k) - h_2(z^{k+1})\Big) \\
&= \Big(\lambda^{k+1} - \lambda^k\Big)^\top \Big(h_2(z^{k+1}) - h_2(z^k)\Big) \;=\; \rho(r_h^{k+1})^\top \Big(h_2(z^{k+1}) - h_2(z^k)\Big),
\end{aligned}$$

where we have used the update rule for λ, see again Line 8 of Algorithm 1. This completes the proof of Lemma 4. □

We are finally ready to prove our main theorem.

Proof of Theorem 1. We know from Lemma 4 that

$$V^k - V^{k+1} \;\geq\; \rho\|r_g^{k+1}\|^2 + \rho\|r_h^{k+1}\|^2 + \rho\|h_2(z^{k+1}) - h_2(z^k)\|^2.$$

Using that $V^k \geq 0$ for every iteration k, see Definition 1, it follows that

$$V^0 \;\geq\; \sum_{k=0}^{\infty} \left(V^k - V^{k+1}\right) \;\geq\; \rho\left(\sum_{k=0}^{\infty} \left(\|r_g^{k+1}\|^2 + \|r_h^{k+1}\|^2 + \|h_2(z^{k+1}) - h_2(z)\|^2\right)\right).$$

The series on the right hand side is absolutely convergent since it holds that $V^0 < \infty$, which follows from the fact that h_2 is an affine function. This absolute convergence implies

$$\lim_{k\to\infty} r_g^k = 0, \quad \lim_{k\to\infty} r_h^k = 0, \quad \text{and} \quad \lim_{k\to\infty} \|h_2(z^{k+1}) - h_2(z^k)\| = 0,$$

i.e., the points x^k and z^k will be primal feasible up to an arbitrarily small error for sufficiently large k. Finally, it follows from Lemmas 2 and 3 that $\lim_{k\to\infty} f^k = f^*$, i.e., the points x^k and z^k are also primal optimal up to an arbitrarily small error for sufficiently large k. □

5 Convex optimization problems with many constraints

Now we are prepared to discuss the main problem that we set out to address in this paper, namely solving general convex optimization problems with many constraints in a distributed setting by distributing the constraints. That is, we want to address optimization problems of the form

$$\begin{array}{ll} \min_x & f(x) \\ \text{s.t.} & g_i(x) \leq 0 \quad i = 1 \dots p \\ & h_i(x) = 0 \quad i = 1 \dots m, \end{array} \tag{8}$$

where $f : \mathbb{R}^n \to \mathbb{R}$ and $g_i : \mathbb{R}^n \to \mathbb{R}^{p_i}$ are convex functions, and $h_i : \mathbb{R}^n \to \mathbb{R}^{m_i}$ are affine functions. In total, we have $p_1 + p_2 + \ldots + p_p$ inequality constraints that are grouped together into p batches and $m_1 + m_2 + \ldots + m_m$ equality constraints that are subdivided into m groups. For distributing the constraints we can assume without loss of generality that $m = p$. That is, we have m batches that each contain p_i inequality and m_i equality constraints.

Again it is easier to work with an equivalent reformulation of Problem 8, where each batch of equality and inequality constraints shares the same variables x_i, namely problems of the form

$$\begin{array}{lll} \min_{x_i,z} & \sum_{i=1}^m f(x_i) & \\ \text{s.t.} & \max\{0,\, g_i(x_i)\} = 0 & i = 1 \dots m \\ & h_i(x_i) = 0 & i = 1 \dots m \\ & x_i = z, & \end{array} \tag{9}$$

where all the variables x_i are coupled through the affine constraints $x_i = z$. To keep our exposition simple, the objective function has been scaled by m in the reformulation.

For specializing our extension of ADMM to instances of Problem 9 we need the Augmented Lagrangian of this problem, which reads as

$$\begin{aligned} L_\rho(x_i, z, \mu_{i,g}, \mu_{i,h}, \lambda) \; = \\ \sum_{i=1}^m f(x_i) \; + \; \frac{\rho}{2} \sum_{i=1}^m \| \max\{0,\, g_i(x_i)\}\|^2 + \sum_i (\mu_{i,g})^\top \left(\max\{0,\, g_i(x_i)\}\right) \\ + \frac{\rho}{2} \sum_{i=1}^m \|h_i(x_i))\|^2 + \sum_i (\mu_{i,h})^\top \left(h_i(x_i)\right) \; + \; \frac{\rho}{2} \sum_{i=1}^m \|x_i - z\|^2 + \sum_i (\lambda_i)^\top (x_i - z), \end{aligned}$$

where $\mu_{i,g}, \mu_{i,h}$, and λ_i are the Lagrange multipliers (dual variables).

Note that the Lagrange function is separable. Hence, the update of the x variables in Line 5 of Algorithm 1 decomposes into the following m independent updates

$$x_i^{k+1} = \operatorname{argmin}_{x_i} \; f(x_i) + \frac{\rho}{2}\| \max\{0,\, g_i(x_i)\}\|^2 + (\mu_{i,g}^k)^\top (\max\{0,\, g_i(x_i)\}) \\ + \frac{\rho}{2}\|h_i(x_i)\|^2 + (\mu_{i,h}^k)^\top (h_i(x_i)) \; + \; \frac{\rho}{2}\|x_i - z^k\|^2 + (\lambda_i^k)^\top (x_i - z^k),$$

that can be solved in parallel once the constraints $g_i(x_i)$ and $h_i(x_i)$ have been distributed on m different, distributed compute nodes. Note that each update is an unconstrained, convex optimization problem, because the functions that need to be minimized are sums of convex functions. The only summand, where this might not be obvious, is $\| \max\{0,\, g_i(x_i)\}\|^2$, but the squared norm of a non-negative, convex function is always convex again.

The update of the z variable in Line 6 of Algorithm 1 amounts to solving the following unconstrained optimization problems

$$z^{k+1} \; = \; \operatorname{argmin}_z \sum_{i=1}^m \frac{\rho}{2}\|x_i^{i+1} - z\|^2 + \sum_{i=1}^m (\lambda_i^k)^\top (x_i^{k+1} - z) \; = \; \frac{\rho \sum_{i=1}^m x_i^{k+1} + \sum_{i=1}^m \lambda_i^k}{\rho \cdot m},$$

and the updates of the dual variables μ_i and λ_i are as follows

$$\mu_{i,g}^{k+1} = \mu_{i,g}^k + \rho \, \max\{0,\, g_i(x_i^{k+1})\}, \quad \mu_{i,h}^{k+1} = \mu_i^k + \rho \, h_i(x_i^{k+1}), \quad \lambda_i^{k+1} = \lambda_i^k + \rho \left(x_i^{k+1} - z^{k+1}\right).$$

That is, in each iteration there are m independent, unconstrained minimization problems that can be solved in parallel on different compute nodes. The solutions of the independent subproblems are then combined on a central node through the update of the z variables and the Lagrange multipliers. Actually, since the Lagrange multipliers $\mu_{i,g}$ and $\mu_{i,h}$ are also local, i.e., involve only the variables x_i^{k+1} for any given index i, they can also be updated in parallel on the same nodes where the x_i^k updates take place. Only the variables z and the Lagrange multipliers λ_i need to be updated centrally.

Looking at the update rules it becomes apparent that Algorithm 1 when applied to instances of Problem 9 is basically a combination of the standard Augmented Lagrangian method [8, 13] for solving convex, constrained optimization problems and ADMM. It combines the ability to solve constrained optimization problems (Augmented Lagrangian) with the ability to solve convex optimization problems distributedly (ADMM).

6 Discussion and conclusions

We have introduced and analyzed an algorithm for solving general convex optimization problems with many constraints in a distributed setting. The algorithm is based on an extension of the alternating direction method of multipliers (ADMM). Potential applications include the smallest enclosing ball problem from computational geometry, where every data point contributes a constraint to the problem, namely the distance of the point to the center of the ball must be at most the radius of the ball. Hence there are many constraints if there are many data points. To the best of our knowledge there exists no other algorithm so far for solving the smallest enclosing ball problem in a distributed setting. Other problems that can involve a large number of constraints are (mixed) packing and covering linear programs. Several parallel and distributed approximation algorithms for these problems have been designed

and analyzed in [1, 5, 9, 10, 15]. Packing and covering linear programs exhibit only linear inequality constraints and thus can be handled by the classical ADMM directly. Indeed, our extension of ADMM becomes the classical ADMM in this special case. Interestingly, ADMM has never been used for packing and covering linear programs before despite the obvious need for distributed solutions. Similarly, to address the generalized matching problem, i.e., a generalization of the standard (weighted) matching problem, where vertices are allowed to be matched more than once, Manshadi et al. [11] have designed a distributed algorithm that can solve mixed packing-covering linear programs, i.e., linear programs that involve packing as well as covering constraints. Again, ADMM has not been used for this type of application before.

Acknowledgments

This work was supported by Deutsche Forschungsgemeinschaft (DFG) under grant GI-711/5-1 within the priority program Algorithms for Big Data.

References

[1] Yair Bartal, John W. Byers, and Danny Raz. Fast, distributed approximation algorithms for positive linear programming with applications to flow control. *SIAM J. Comput.*, 33(6):1261–1279, 2004.

[2] Stephen P. Boyd, Neal Parikh, Eric Chu, Borja Peleato, and Jonathan Eckstein. Distributed optimization and statistical learning via the alternating direction method of multipliers. *Foundations and Trends in Machine Learning*, 3(1):1–122, 2011.

[3] Kaspar Fischer, Bernd Gärtner, and Martin Kutz. Fast smallest-enclosing-ball computation in high dimensions. In *11th Annual European Symposium on Algorithms (ESA)*, pages 630–641. Springer, 2003.

[4] Daniel Gabay and Bertrand Mercier. A dual algorithm for the solution of nonlinear variational problems via finite element approximation. *Computers & Mathematics with Applications*, 2(1):17 – 40, 1976.

[5] Naveen Garg and Jochen Könemann. Faster and simpler algorithms for multicommodity flow and other fractional packing problems. In *39th Annual Symposium on Foundations of Computer Science, FOCS*, pages 300–309. IEEE Computer Society, 1998.

[6] Bernd Gärtner. Fast and robust smallest enclosing balls. In *7th Annual European Symposium on Algorithms (ESA)*, pages 325–338. Springer, 1999.

[7] R. Glowinski and A. Marroco. Sur l'approximation, par lments finis d'ordre un, et la rsolution, par pnalisation-dualit d'une classe de problmes de dirichlet non linaires. *ESAIM: Mathematical Modelling and Numerical Analysis - Modlisation Mathmatique et Analyse Numrique*, 9(R2):41–76, 1975.

[8] Magnus R. Hestenes. Multiplier and gradient methods. *Journal of Optimization Theory and Applications*, 4(5):303–320, 1969.

[9] Slobodan Jelic, Sören Laue, Domagoj Matijevic, and Patrick Wijerama. A fast parallel implementation of a PTAS for fractional packing and covering linear programs. *International Journal of Parallel Programming*, 43(5):840–875, 2015.

[10] Christos Koufogiannakis and Neal E. Young. Distributed algorithms for covering, packing and maximum weighted matching. *Distributed Computing*, 24(1):45–63, 2011.

[11] Faraz Makari Manshadi, Baruch Awerbuch, Rainer Gemula, Rohit Khandekar, Julián Mestre, and Mauro Sozio. A distributed algorithm for large-scale generalized matching. *PVLDB*, 6(9):613–624, 2013.

[12] Damon Mosk-Aoyama, Tim Roughgarden, and Devavrat Shah. Fully distributed algorithms for convex optimization problems. *SIAM Journal on Optimization*, 20(6):3260–3279, 2010.

[13] M. J. D. Powell. Algorithms for nonlinear constraints that use lagrangian functions. *Mathematical Programming*, 14(1):224–248, 1969.

[14] E. Alper Yildirim. Two algorithms for the minimum enclosing ball problem. *SIAM Journal on Optimization*, 19(3):1368–1391, 2008.

[15] Neal E. Young. Sequential and parallel algorithms for mixed packing and covering. In *42nd Annual Symposium on Foundations of Computer Science, FOCS*, pages 538–546, 2001.

[16] Minghui Zhu and Sonia Martínez. On distributed convex optimization under inequality and equality constraints. *IEEE Trans. Automat. Contr.*, 57(1):151–164, 2012.